7

Biology Meets Economics: Evolutionary Stability and the Origins of Dynamic Game Theory

Il n'y a qu'un pas entre le sublime et le ridicule.
Leo Tolstoy

7.1 The Origins of Evolutionary Stability

In 1973 the biologist John Maynard Smith and the mathematician G. R. Price wrote an article in *Nature* showing how game theory applies to the behavior of animals (§7.10). Maynard Smith went on to write a book on the subject (Maynard Smith 1982), which has become a classic. The idea of applying game theory to animals—and not just the higher primates, but fish, dung beetles, fireflies, and pond scum as well—seemed strange at the time, since game theory had always been the preserve of hyperrationality. Animals hardly fit the bill. Maynard Smith made three critical shifts—one in the concept of a strategy, the second in the concept of equilibrium, and a third in the nature of agent interactions.

Strategy: In classical game theory, *players* have strategy sets from which they choose particular strategies. In biology, *species* have strategy sets (genotypic variants), of which *individuals* inherit one or another variant (perhaps mutated), which they then play in their strategic interactions. This extends nicely to the treatment of culture in human society. We say that *society* has the strategy set (the set of alternative cultural forms), and *individuals* inherit or choose among them.

Equilibrium: In place of Nash equilibrium, Maynard Smith and Price used the *evolutionary stable strategy* concept. A strategy is evolutionarily stable if a whole population using that strategy cannot be invaded by small set of mutant genotypes. Similarly, a cultural form is evolutionarily stable if, upon being adopted by all members of a society (firm, family,

VISIT…

etc.), no small group of individuals using an alternative cultural form can invade. We thus move from explaining the actions of individuals to modeling the diffusion of forms of behavior ('strategies') in society.

Player interactions: In place of the one-shot and repeated games of classical game theory, Maynard Smith introduced the notion of the *repeated, random pairing of agents* who play strategies based on their genome but not the previous history of play.

Evolutionary stability is particularly useful because it says something about the dynamic properties of a system without being committed to any particular dynamic model. As we shall see, however, the precise relationship with dynamic stability is by no means straightforward (§9.7).

7.2 Properties of Evolutionary Stable Strategies

Consider a two-player normal form game in which both players have the set $S = \{s_1, \ldots, s_n\}$ of pure strategies, and the payoffs to an agent playing $s_i \in S$ and another agent playing $s_j \in S$ are π_{ij} for the first and π_{ji} for the second. We call such a game *symmetric in payoffs*. In addition, we assume the agents cannot condition their play on whether they are 'player one' or 'player two.' We call such a game *symmetric in strategies*. If a game is symmetric in both payoffs and strategies, we simply call the game *symmetric*. We call $A = (\pi_{ij})$ the *matrix* of the symmetric game (notice that A represents only the payoffs for the row player, since the payoffs to the column player are just the transpose of A).

Let $\mathcal{G}$ be a symmetric game with matrix A (we'll call it the *stage game*), and large population of agents. In each period $t = 1, 2, \ldots$, agents are randomly paired and they play the stage game $\mathcal{G}$. Each agent is of type i for some $s_i \in S$, meaning that the agent uses strategy s_i in the stage game. If the proportion of agents of type j is p_j at a particular time, we say the *state* of the population is $\sigma = p_1 s_1 + \ldots + p_n s_n$. The payoff at that time to a player of type i when the state of the population is σ, is defined by

$$\pi_{i\sigma} = \sum_{j=1}^{n} \pi_{ij} p_j,$$

which is the player's expected payoff before being assigned a particular partner. These conditions define a new game, called the *evolutionary game* corresponding to the stage game $\mathcal{G}$. In some evolutionary games individuals can play mixed strategies, in which case an agent who plays

a mixed strategy $\tau = q_1 s_1 + \ldots + q_n s_n$ is called of 'type τ', and has payoff

$$\pi_{\tau\sigma} = \sum_{i,j=1}^{n} q_i \pi_{ij} p_j,$$

when the state of the population is σ.

Suppose the state of the population is $\sigma = p_1 s_1 + \ldots + p_n s_n$. The expected payoff to a randomly chosen member of the population is thus just $\pi_{\sigma\sigma}$. If we replace a fraction $\epsilon > 0$ of the population with a 'mutant' of type j, the new state of the population is

$$\tau = (1 - \epsilon)\sigma + \epsilon s_j,$$

so the payoff to a randomly chosen nonmutant is

$$\pi_{\sigma\tau} = (1 - \epsilon)\pi_{\sigma\sigma} + \epsilon \pi_{\sigma j}$$

and the payoff to a mutant is

$$\pi_{j\tau} = (1 - \epsilon)\pi_{j\sigma} + \epsilon \pi_{jj}$$

We say the mutant type can *invade* the population if $\sigma \neq \tau$ and for sufficiently small $\epsilon > 0$,

$$\pi_{j\tau} \geq \pi_{\sigma\tau},$$

which says that a mutant does at least as well against the new population as does a nonmutant. We say σ is *evolutionarily stable* if it cannot be invaded by any mutant type.

A Nash equilibrium in an evolutionary game can consist of a *monomorphic* population of agents, each playing the same mixed strategy, or as a *polymorphic* population, a fraction of the population playing each of the underlying pure strategies in proportion to its contribution to the mixed Nash strategy. The two interpretations are interchangeable under many conditions, and we shall not commit ourselves exclusively to either interpretation. We will however assume that mutants can employ mixed strategies in applying the evolutionary stability criterion, because as we shall see later (§9.6), with this assumption evolutionarily stable strategies have powerful dynamic properties that resistance to invasion by pure-strategy mutants do not.

The following exercise asks you to derive some of the more important properties of evolutionarily stable strategies:

(a) Show that σ is evolutionarily stable if and only if, for any mutant type τ, we have

$$\pi_{\sigma\sigma} \geq \pi_{\tau\sigma},$$

and if $\pi_{\sigma\sigma} = \pi_{\tau\sigma}$, then

$$\pi_{\sigma\tau} > \pi_{\tau\tau}.$$

This says that σ is evolutionarily stable if and only if a mutant cannot do better against an incumbent than an incumbent can do against another incumbent, and if a mutant does as well as an incumbent against another incumbent, then an incumbent must do better against a mutant than a mutant does against another mutant. Notice here that we are assuming mutants can used mixed strategies.

(b) Show that an evolutionary stable strategy is a Nash equilibrium.

(c) Show that every strict Nash equilibrium (§4.4) in an evolutionary game is evolutionarily stable.

(d) Suppose the stage game has two pure strategies. Show that if $\pi_{11} \neq \pi_{21}$ and $\pi_{12} \neq \pi_{22}$, then the game has an evolutionarily stable strategy.

(e) Using the same notation, show that the stage game has a strictly mixed Nash equilibrium if and only if $\pi_{11} > \pi_{21}$ and $\pi_{22} > \pi_{12}$, or $\pi_{11} < \pi_{21}$ and $\pi_{22} < \pi_{12}$. The equilibrium is evolutionarily stable only if the second set of inequalities holds.

(f) *Show that an evolutionary game whose stage game has a finite number of pure strategies can have only a finite number of evolutionarily stable strategies.

(g) **Show that σ is evolutionarily stable if and only if there is some $\epsilon > 0$ such that $\pi_{\sigma\tau} > \pi_{\tau\tau}$ for all τ within distance ϵ of σ.[1] This implies that an evolutionarily stable strategy σ is an *isolated Nash equilibrium*, in the sense that there is an $\epsilon > 0$ such that no strategy $\tau \neq \sigma$ within distance ϵ of σ is a Nash equilibrium. It also clearly proves (f).

[1]By the *distance* between two strategies $\sigma = \sum_i p_i s_i$ and $\tau = \sum_i q_i s_i$ is the distance in $\mathbf{R}^n$ between the points $(p_1, \ldots, p_n)$ and $(q_1, \ldots, q_n)$. If your are stumped, see and Hofbauer and Sigmund 1998.

(h) Show that if $\sigma \in S$ is a completely mixed evolutionarily stable strategy (i.e., it uses all pure strategies with positive probability), then it is the unique Nash equilibrium of the game, and the condition

$$\pi_{\sigma\tau} > \pi_{\tau\tau}$$

holds for all mixed strategies τ, $\tau \neq \sigma$.

7.3 When are Evolutionarily Stable Strategies Unbeatable?

It is easy to show that x is evolutionarily stable in the game to the right. We shall see later that it is also asymptotically stable in the replicator dynamic (§9.7). Nevertheless it is not an *unbeatable strategy*, in the sense of always having the highest payoff when invaded by multiple mutants.

	x	y	z
x	1,1	1,1	0,0
y	1,1	0,0	1,0
z	0,0	0,1	0,0

(a) Show that x is evolutionarily stable;

(b) Show that if a fraction ϵ_y of y-players and a fraction $\epsilon_z > \epsilon_y$ of z-players simultaneously invade, then y has a higher payoff than x.

To complicate the picture, some game theorists have *defined* evolutionary stability as 'unbeatability' in the above sense. In a famous article, **?**) showed that "No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game," and Farrell and Ware (1989) extended this by showing that no mixed strategy using a finite number of pure strategies is evolutionarily stable. Finally, Lorberbaum (1994) extended this to all non-deterministic strategies, and Bendor and Swistak (1995) showed that for a sufficiently low discount rate, no pure strategy is evolutionarily stable in any nontrivial repeat game. In all cases, however, evolutionarily stability is interpreted as 'unbeatability' in the above sense. But unbeatability is not a very important concept, since as far as I know, it has no interesting dynamic properties.

7.4 Cooperative Fishing

In a certain fishing village, two fisherman gain from having the nets put out in the evening. However, the fishermen benefit equally whether or

not they share the costs of putting out the nets. Suppose the expected catch is ν, and the cost of putting out the nets to each is c_1 if he does it alone, and the cost to each is $c_2 < c_1$ if they do it together. We assume $\nu/2 > c_1$, so it is worth while for a fisherman to put out the nets even if he has to do it alone. But since $c_2 < c_1$, he prefers help. On the other hand, by free-riding on the first fisherman's effort (i.e., by not helping), the other fisherman gets $\nu/2$ anyway.

What is the normal form game, where each fisherman has the available strategies Put Out (P) and Free Ride (F)? Show that there is a unique mixed strategy equilibrium and that this strategy is evolutionarily stable.

7.5 A Nash Equilibrium that is not Evolutionarily Stable

Suppose agents consume each other's products, but not their own. An agent can produce one or two units per period at cost 1, and then meets another consumer. They can agree to exchange either one or two units. The utility of consumption is 2 per unit consumed. The first strategy is thus "exchange equal for equal, but at most one unit of the good," and strategy two is "exchange equal for equal, but at most two units of the good." The payoff matrix is shown at the right. Show that one of the Nash equilibria consists of evolutionarily stable strategies, but the other does not. What does this say about evolutionary stability and the elimination of weakly dominated strategies?

	1	2
1	$1,1$	$1,1$
2	$1,1$	$2,2$

7.6 Rock, Paper, and Scissors is not Evolutionarily Stable

A Nash equilibrium that is not evolutionarily stable may nevertheless be quite important. Consider, for instance, Rock, Paper, and Scissors (§4.15). Show that the unique, completely mixed Nash equilibrium to this game is not evolutionarily stable. We will see later (§9.11) that under a replicator dynamic, Rock, Paper, and Scissors traces out closed orbits around the equilibrium (1/3,1/3,1/3), as suggested in the example of the lizard *Uta stansburiana* (§4.15).

7.7 Sex Ratios as Evolutionary Stable Strategies

The sex ratio game in §4.16 is not an evolutionary game, strictly speaking, since females do not meet pairwise and play a stage game. But it has a payoff structure like an evolutionary game, since $f(u, v)$ is the payoff to a female using sex ratio u when the population is using v. Show that in this sense, the strategies used in the mixed strategy Nash equilibrium found in the first part of §4.16 are evolutionarily stable.

7.8 Invasion of the Pure Strategy Mutants

1,1	1,1	1,1
1,1	0,0	a,a
1,1	a,a	0,0

It is possible that a Nash equilibrium be immune to invasion by any pure strategy mutant, but not by an appropriate *mixed* strategy mutant. This is the case with respect to the game at the right, assuming $a > 2$. Here the first strategy is evolutionarily stable if only pure strategy mutants are allowed, but not if mixed strategy mutants are allowed.

Show that a mixed strategy using pure strategies two and three each with probability 1/2 can invade a Nash equilibrium consisting of strategy 1 alone. Is there an evolutionarily stable strategy using pure strategies two and three? Since mutants are normally considered to be rare, it is often plausible to restrict consideration to single mutant types, or to mixed strategies that include only pure strategies used in the Nash equilibrium, plus at most one mutant.

7.9 Multiple Evolutionary Stable Strategies

5,5	7,8	2,1
8,7	6,6	5,8
1,2	8,5	4,4

Using the matrix to the right, show that there are two evolutionarily stable strategies, one using the first two rows and columns, and the second using the second and third strategies. Show that there is also a completely mixed Nash equilibrium that is stable against invasion by pure strategies, but not evolutionarily stable.

Prove the latter either by finding a mixed strategy that does invade (Hint: try one of the evolutionarily stable strategies) or use a previously proved theorem. If you want to cheat, look up a paper by Haigh (1975), which develops a simple algorithm for determining whether a Nash equilibrium is evolutionarily stable.

7.10 The Logic of Animal Conflict

In many of species throughout the animal kingdom, males compete for access to females. Since the result of this competition is reproductive success, it is not surprising that males in many species are bigger and strong than females, and often have awesome fighting and reproductive equipment (Alcock 1993). Is it not surprising, then that so few animals have developed truly *deadly* instruments to use against their compatriots? Venomous snakes wrestle, but do not bite each other in combat. Deer have impressive antlers and butt head furiously, but their antlers are not lethal weapons, and they rarely kill each other. Horns in many species point backwards, or are so tightly curled up that they cannot inflict much damage. Even animals that have very dangerous equipment (e.g., the chimpanzee's bite or the elephant's tusks) deploy it rarely.

The common view among animal behaviorists, brilliantly developed in Wynne-Edwards 1962 book *Animal Dispersion in Relation to Social Behavior* , was that limitations on intraspecies warfare evolved because it was "for the good of the species." This explanation never sat well with many biologists, however, because it is not clear how the "good of the species" becomes encoded in the genes of individuals. Provoked by Wynne-Edwards, George C. Williams published *Adaptation and Natural Selection* in 1966, a devastating attack on the notion of *group selection* implicit in Wynne-Edwards' argument. The fact that a species exists, Williams noted, of course means that the males of the species do not kill each other off. But we cannot use the existence of the species as an *explanation* of why they don't kill each other off.

There were two types of response to Williams' challenge. One was to

develop analytically rigorous models of group selection, specify the conditions under which they are operative, and argue that these conditions are biologically plausible. We touch on this body of research in §11.7. The more influential response, however, was to assert that what appears to be altruistic behavior is in fact self-interest in disguise. Prominent among those holding this position were Maynard Smith and Price, whose "The Logic of Animal Conflict" x(1973) gave birth to evolutionary game theory, and long before Axelrod and Hamilton's influential studies (Axelrod and Hamilton 1981, Axelrod 1984), used computer simulations to support their theoretical arguments.[2]

In Maynard Smith and Price's model, pairs of individuals meet randomly and play a stage game repeated for a fixed number of rounds. Players have three possible actions, 'dangerous' (D), 'conventional' (C), and 'retreat' (R). By playing D, the opponent receives serious injury (payoff -100) with 10% probability, and minor injury (payoff -2) with probability 90%. A player whose opponent retreats (R) has a payoff of 60, and the stage game is over. The 'Mouse' strategy is to play C until seriously injured, then retreat. The 'Hawk' strategy is to play D until seriously injured, then retreat. The 'Bully' strategy is to play D until the opponent plays D twice, then retreat. The 'Retaliator' strategy is to play C, then do what your opponent did on the last move (Tit-for-Tat). The 'Prober-Retaliator' strategy is to act like the Retaliator, except that with some positive probability, play D, and continue to play D unless the opponent retaliates by playing D. They came up with the payoff matrix shown in Figure 7.1. Since the game is symmetric, only the payoffs for the row player are shown.

To find the Nash equilibria and evolutionarily stable strategies in the game determined by this payoff matrix, I used computer software that I wrote to solve two-player normal form games. You should check the following:

(a) There is a symmetric Nash equilibrium that is immune to invasion by pure strategies in which the population consists of 57.6% Hawk

[2]Actually, Maynard Smith and Price did not explicitly use game theory in this paper though they employ strategies and a payoff matrix, and did not define an evolutionarily stable strategy appropriately. Moreover, they did not investigate the equilibria of their model in any systematic way, and made incorrect inferences from their own model. Nevertheless, this paper is a brilliant classic. "What can be said," said Wittgenstein, "can be said clearly." But the *first time* something is said, it is rarely said clearly, and researchers who insist on clarity at all costs rarely make important discoveries. Maynard Smith and Price's paper is a wonderful case in point.

	Mouse	Hawk	Bully	Retaliator	Prober-Retaliator
Mouse	29.0	19.5	19.5	29.0	17.2
Hawk	80.0	-19.5	74.6	-18.1	-18.9
Bully	80.0	4.9	41.5	11.9	11.2
Retaliator	29.0	-22.3	57.1	29.0	23.1
Prober-Retaliator	56.7	-20.1	59.4	26.9	21.9

Figure 7.1: Average Payoffs in Simulated Intraspecific Contests

and 42.4% Bully. This equilibrium is in fact evolutionarily stable, but that is quite hard to check. Using the techniques of Chapters 8 and 9, we can show that this equilibrium is asymptotically stable in the replicator dynamic, which implies that it is evolutionarily stable. Alternatively, you can use the technique described in Haigh (1975).

(b) There is a symmetric Nash equilibrium that is not evolutionarily stable, with 38.9% Hawk, 44.1% Bully, and 17.0% Prober-Retaliator.

(c) (Retaliate,Retaliate) is a symmetric Nash equilibrium, but it is not evolutionarily stable, since it can be invaded by Mouse;

(d) (Mouse,Hawk) is a Nash equilibrium, but since the game is symmetric and the players cannot tell whether they are row or column players, this stage game equilibrium cannot occur in the evolutionary game. There is another such asymmetric equilibrium of the stage game, in which one player uses (Hawk, Bully, Prober-Retaliator) with probabilities (0.366,0.327,0.308) and the other player uses (Mouse,Bully,Prober-Retaliator) with probabilities (0.377,0.296,0.326);

Notice that there is a *unique* evolutionarily stable mixed strategy in this game, consisting of Hawks and Bullies. Maynard Smith and Price do not find this evolutionarily stable strategy, and speculate that the species would evolve into mainly Retaliators and Prober-Retaliators, which is not the case.

7.11 Hawks, Doves and Bourgeois

	H	D
H	$(v-w)/2$	v
D	0	$v/2-t$

Consider a population of birds that fight over valuable territory. There are two possible strategies. The 'hawk' (H) strategy is to escalate battle until injured or your opponent retreats. The 'dove' (D) strategy is to display hostility, but retreat before sustaining injury if your opponent escalates. The payoff matrix given in the figure to the right, where v is the value of territory, w is the cost of injury, and t is the cost of protracted display when D plays against D (in this and the next figure we have included only the payoff to player 1, since the payoff to player 2 is the payoff to player 1 symmetrically across the diagonal). The birds can play mixed strategies, but they are physically indistinguishable.

(a) Justify the entries in the payoff matrix;

(b) Show that if $w > v$, neither H nor D is an evolutionarily stable strategy;

(c) Suppose the birds all play $\alpha H + (1-\alpha)D$. Show that if $t > 0$ and $w > v$ there is a unique $\alpha \in (0,1)$ such that this mixed strategy is evolutionarily stable;

(d) Suppose a new mutant comes along, a 'bourgeois' (B) who plays H if it was the first to occupy the territory, and plays D if it comes upon another bird's territory. Show that, under plausible conditions, when we add B to the birds' repertoire, we get the payoff matrix in Figure 7.2. What relations among v, w, and t must obtain in order that B be evolutionarily stable?

7.12 Trogs and Farfel

Trogs float in a primordial ooze. They make farfel, which is the only food they eat. But a trog cannot eat its own farfel. Rather, each trog secretes a unit of farfel, then floats around in the ooze until it encounters another trog, after which they exchange their farfel. A trog can make either Good farfel or Bad farfel. It costs a trog x ergs to make a unit of bad farfel and y ergs to make a unit of good farfel ($y > x$), and the energy derived from consuming good farfel is z and from consuming bad

	H	D	B
H	$(v-w)/2$	v	$3v/4-w/4$
D	0	$v/2-t$	$v/4-t/2$
B	$(v-w)/4$	$3v/4-t/2$	$v/2$

Figure 7.2: Hawks, Doves, and Bourgeois

farfel is w ($z > w$). A Trog's payoff from a transaction is the net number of ergs obtained.

(a) Show that the trogs face a prisoners' dilemma if and only if $z-w > y-x$.

The trogs can recognize each other, so each encounter between trogs is one in a series of such encounters, and thus is part of a repeated prisoner's dilemma between the two trogs. Trogs are genetically wired to play strategies in the normal form of the repeated prisoner's dilemma.

Tit-for-Tat is a strategy defined as follows. The first time you meet a new trog, give it Good farfel. Each time you meet this particular trog again, give the trog the same type of farfel it gave you in your previous encounter (§5.11).

Trogs use a discount factor δ, with $0 < \delta < 1$, to add up energy over time: an erg in the next period is worth δ ergs today, because the trog might die before the next period arrives. Assume the expected length of time before a trog meets its current partner again is one time period.

(b) Show that Tit-for-Tat is not evolutionarily stable, and can be invaded by the strategy "always give good farfel."

We say a particular strategy is *weakly evolutionarily stable* if, when all Trogs follow this strategy no mutant following a different strategy can invade it in the sense of doing better than the incumbent trogs in exchanges.

(c) Show that Tit-for-Tat is weakly evolutionarily stable if and only if

$$\delta > \frac{y-x}{z-w}.$$

Hint: Let G and B represent "give Good farfel" and "give Bad farfel," respectively. First show that playing B forever against Tit-for-Tat is not a best response if and only if the above condition holds. Then show that playing B for a certain number of rounds, followed by playing G, is not a best response if and only if the above condition holds.

(d) Show that "always give Bad farfel" is evolutionarily stable.

(e) If you are musically literate, compose and sing the song "You Give Good Farfel."

7.13 Evolutionary Stability in Finite Populations

Consider a population in which agents are randomly paired in each period and each pair plays a 2×2 game. Let $r_{\mu\nu}$ be the payoff to playing μ when your partner plays ν. Suppose agents are 'hardwired' to play a particular strategy, and let $r(\mu)$ and $r(\nu)$ be the expected payoffs to an μ-type and a ν-type agent, respectively.

Suppose there are n agents, m of which play the 'mutant' strategy μ, the rest playing the 'incumbent' strategy ν. Then we have

$$r(\mu) = \left(1 - \frac{m-1}{n}\right) r_{\mu\nu} + \frac{m-1}{n} r_{\mu\mu}$$
$$r(\nu) = \left(1 - \frac{m}{n}\right) r_{\nu\nu} + \frac{m}{n} r_{\nu\mu}.$$

It follows that

$$r(\nu) - r(\mu) = \left(1 - \frac{m}{n}\right)(r_{\nu\nu} - r_{\mu\nu}) + \frac{m}{n}(r_{\nu\mu} - r_{\mu\mu}) + \frac{1}{n}(r_{\mu\mu} - r_{\mu\nu}). \quad (7.1)$$

We say a strategy ν is *noninvadible* if there is an $\epsilon > 0$ such that for all feasible mutants $\mu \neq \nu$ and all positive m such that $m/n < \epsilon$, $r(\nu) > r(\mu)$. When this condition fails, we say ν is *invadible*. We say a strategy ν is *Nash* if ν is a best reply to itself, or equivalently, if there is a Nash equilibrium in which only ν is played. We say a strategy ν is *evolutionarily stable* if there is a population size n such ν is noninvadible for all populations of size n or greater.

(a) Show that it is possible for a non-Nash strategy to be noninvadible, even by a Nash strategy.

(b) Show that is possible for a Nash strategy to be invadible.

(c) Show that strategy ν is evolutionarily stable if and only if ν is a Nash strategy, and for any μ that is a best reply to ν, ν is a better reply to μ than μ is to itself, or if μ is as good a reply to itself as ν is to μ, then μ is a better reply to itself than μ is to ν. Equivalently, for any mutant μ, either $r_{\nu\nu} > r_{\mu\nu}$, or $r_{\nu\nu} = r_{\mu\nu}$ and $r_{\nu\mu} > r_{\mu\mu}$, or $r_{\nu\nu} = r_{\mu\nu}$ and $r_{\nu\mu} = r_{\mu\mu}$ and $r_{\mu\mu} > r_{\mu\nu}$.

(d) The above conditions are not those of Maynard Smith, which state that ν is evolutionarily stable if and only if ν is Nash and for any μ that is a best reply to n, ν is a better reply to μ than μ is to itself, or equivalently, for any mutant μ, either $r_{\nu\nu} > r_{\mu\nu}$, or $r_{\nu\nu} = r_{\mu\nu}$ and $r_{\nu\mu} > r_{\mu\mu}$.

Show that we can derive Maynard Smith's conditions by letting $m, n \to \infty$ in (7.1) in such a manner that $m/n = \epsilon$, but the limit argument cannot be use to conclude that $r(\nu) > r(\mu)$ in the 'large finite' case.

(e) Let $p = m/n$, $a = r_{\nu\nu} - r_{\mu\nu}$, $b = r_{\nu\mu} - r_{\mu\mu} - r_{\nu\nu} + r_{\mu\nu}$, and $c = r_{\mu\mu} - r_{\mu\nu}$. Then (7.1) becomes

$$r(\nu) - r(\mu) = \frac{1}{n}(na + mb + c). \tag{7.2}$$

Suppose ν can be invaded by mutant strategy μ, and the system follows any dynamic in which a strategy with a higher payoff increases in frequency. Show that if $n(a+b) > c$, μ will invade until μ is the largest integer less than $-(na+c)/b$. Otherwise μ will invade until ν is extinct.

In the case of partial invasion in the last case, we say μ is *quasi-evolutionarily stable*. Note that μ is quasi-evolutionary stable with respect to ν for very large n if and only if μ and ν are part of a completely mixed Nash equilibrium (assuming there are no other feasible pure strategies).

7.14 Evolutionary Stability in Asymmetric Games

Many situations can be modeled as evolutionary games, except for the fact that the two players are not interchangeable. For instance, in One-card, Two-round Poker with Bluffing (§4.11), the player going first has a

completely different set of strategies from the player going second. Yet, despite the lack of symmetry, we simulated the game quite nicely as an evolutionary game. Analogous situations include interactions between predator and prey, boss and worker, male and female, incumbent and intruder, among a host of others.

This is not simply a technicality—it makes no sense to say that a "mutant meets its own type" when the game is asymmetric, so evolutionary stability has no meaning. The obvious way around this problem is to define a homogeneous set of players who in each period are paired randomly, one of the pair being randomly assigned to be player 1, and the other to be player 2. We may call this the 'symmetric version' of the asymmetric evolutionary game. However, *an evolutionarily stable strategy in the symmetric version of an asymmetric evolutionary game must be a strict Nash equilibrium*—that is, each type in the asymmetric game must use exactly *one* pure strategy (Selten 1980). To see this, suppose there is a Nash equilibrium $u = (\sigma_1, t_2)$ of the symmetric version, where a player uses strictly mixed strategy σ_1 when assigned to be player 1, and uses t_2 (pure or mixed) when assigned to be player 2. Consider a mutant that uses $v = (s_1, t_2)$, where s_1 is a pure strategy that appears with positive weight in σ_1. Then v does as well against u as u does against itself, and v does as well against v as u does against v. All this is true because the payoff to s_1 against t_2 in the asymmetric game is equal to the payoff to σ_1 against t_2 by the Fundamental Theorem (§4.3).

It follows that *mixed strategy Nash equilibria in asymmetric evolutionary games are never evolutionarily stable in the symmetric version of the game.* As we shall see later, this situation reflects the fact that mixed strategy Nash equilibria in asymmetric evolutionary games with a replicator dynamic are never asymptotically stable (§9.14). Some game theorists consider this a weakness of evolutionary game theory (Mailath 1998), but in fact it reflects a deep and important regularity of social interaction—in asymmetric evolutionary games, the frequency of different types of behavior goes through periodic cycles through time.

For a dramatic example of this important insight, we return to our simulation of One-Card Two-Round Bluff (Figure 4.2). In this simulation I lowered the mutation rate to 1%, and ran the simulation ran for 300,000 rounds. The results are shown in Figure 7.3 for one of the player 1 types, *rrbb* (bluff all the way). Note that the *average frequency* of each strategy settles down to the theoretically predicted equilibrium value, but the *round-to-round frequencies* fluctuate wildly in the medium run. Strategy *rrbb*, which has the equilibrium frequency of about 50%, sometimes goes

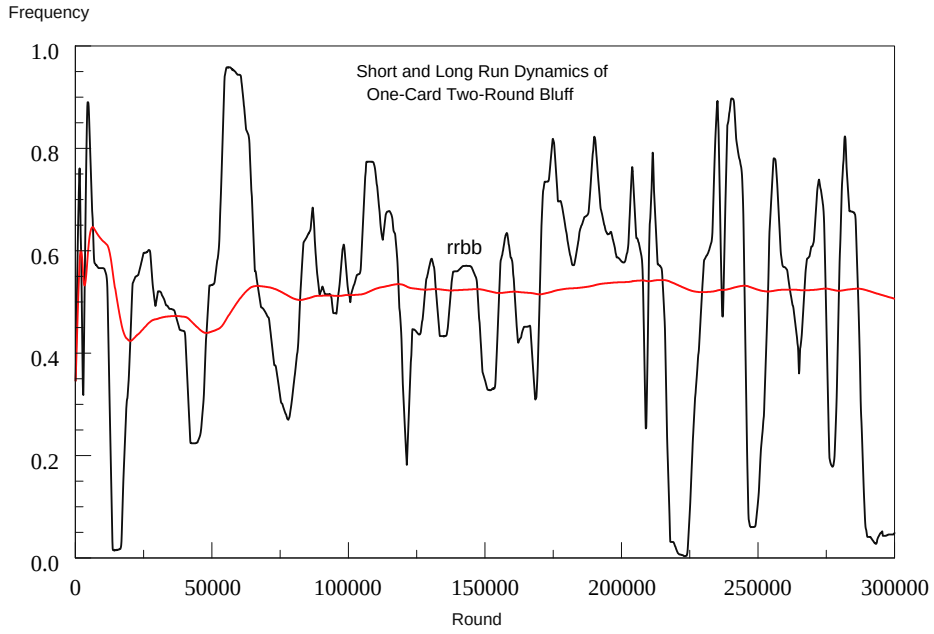

Figure 7.3: Simulation of One-Card Two-Round Bluff

for thousands of rounds with frequency under 5% or over 90%.